arXiv:1702.08896v2 [stat.ML] 13 Mar 2017

Deep and Hierarchical Implicit Models

Dustin Tran [1] **Rajesh Ranganath** [2] **David M. Blei** [1]

Abstract

Implicit probabilistic models are a flexible class for modeling data. They define a process to simulate observations, and unlike traditional models, they do not require a tractable likelihood function. In this paper, we develop two families of models: *hierarchical implicit models* and *deep implicit models*. They combine the idea of implicit densities with hierarchical Bayesian modeling and deep neural networks. The use of implicit models with Bayesian analysis has been limited by our ability to perform accurate and scalable inference. We develop *likelihood-free variational inference (*LFVI*)*. Key to LFVI is specifying a variational family that is also implicit. This matches the model's flexibility and allows for accurate approximation of the posterior. Our work scales up implicit models to sizes previously not possible and advances their modeling design. We demonstrate diverse applications: a large-scale physical simulator for predator-prey populations in ecology; a Bayesian generative adversarial network for discrete data; and a deep implicit model for text generation.

1. Introduction

Consider a model of coin tosses. With probabilistic models, one typically posits a latent probability, and supposes each toss is a Bernoulli outcome given this probability (Murphy, 2012; Gelman et al., 2013). After observing a collection of coin tosses, Bayesian analysis lets us describe our inferences about the probability.

However, we know from the laws of physics that the outcome of a coin toss is fully determined by its initial conditions (say, the impulse and angle of flip) (Keller, 1986; Diaconis et al., 2007). Therefore a coin toss' randomness does not originate from a latent probability but in noisy initial parameters. This alternative model incorporates the physical system, better capturing the generative process. Furthermore the model is *implicit*, also known as a simulator: we can sample data from its generative process, but we may not have access to calculate its density (Diggle & Gratton, 1984; Hartig et al., 2011).

Implicit probabilistic models have a rich history, dating back to modeling philosophies since Laplace's demon (Laplace, 1814) and fiducial probability (Fisher, 1930). They argue for all generative models to generate data from a deterministic equation given parameters and random noise (i.e., without a "likelihood"). Such models pervade physical systems in fields such as population genetics, statistical physics, and ecology (Pritchard et al., 1999; Anelli et al., 2008; Beaumont, 2010). They also underlie structural equation models in causality (Pearl, 2000).

In this paper, we introduce *hierarchical implicit models*, a class of Bayesian hierarchical models which only assume a process that generates samples. This class encompasses most simulators in the literature, including those employed in generative adversarial networks (GANs) (Goodfellow et al., 2014), an approach which uses neural networks to transform noise into data samples. By introducing sample generation processes at multiple layers of latent variables, we also develop *deep implicit models*, a class of models that lets us discover not only rich data structure but also rich latent structure. Deep implicit models generalize deep generative models such as sigmoid belief nets (Neal, 1990) and deep latent Gaussian models (Rezende et al., 2014).

Previous work with implicit models has been limited by our ability to perform accurate and scalable Bayesian inference. There is no likelihood available (e.g., in a complex simulator) or it is intractable. Inspired by ideas from ratio estimation (Sugiyama et al., 2012; Mohamed & Lakshminarayanan, 2016), we develop *likelihood-free variational inference (*LFVI*)*. Variational inference posits a family of distributions over the latent variables and then optimizes to find the member closest to the posterior (Jordan et al., 1999). Key to LFVI is specifying a variational family that is also implicit. This matches the model's flexibility and allows for accurate inferences. Further, we sketch a uniqueness proposition: under certain conditions, the KL divergence is the only tractable objective for variational optimization.

[1]Columbia University, New York, NY, USA [2]Princeton University, Princeton, NJ, USA. Correspondence to: Dustin Tran <dustin@cs.columbia.edu>.

This work has diverse applications. First, we analyze a classical problem from the approximate Bayesian computation (ABC) literature, where the model simulates an ecological system (Beaumont, 2010). We use a large-scale data set not possible with traditional methods. Second, we analyze a Bayesian GAN, which is a GAN with a prior over its weights. This allows GANs to quantify uncertainty and improve data efficiency. Moreover, we apply it to discrete data; this setting isn't possible with traditional estimation algorithms for GANs. Third, we analyze deep implicit models for text with recurrent neural networks. Our work enables implicit models at scale and advances their modeling design.

Related Work. The typical method for dealing with implicit models is ABC (Beaumont, 2010; Marin et al., 2012). ABC steps around the intractable likelihood by measuring the closeness of samples from the simulator to real observations via summary statistics. While successful in many domains, ABC has many shortcomings. First, the results generated by ABC depend heavily on the chosen summary statistics and the closeness measure. Second, as the dimensionality grows, closeness becomes harder to achieve. This is the classic curse of dimensionality.

GANs have seen much interest since their conception (Goodfellow et al., 2014), providing an efficient method for estimation in neural network simulator models. Larsen et al. (2016) propose a hybrid of variational methods and GANs for improved reconstruction. Chen et al. (2016) apply information penalties to disentangle factors of variation. Donahue et al. (2017); Dumoulin et al. (2017) propose to match on an augmented space, simultaneously training the model and an inverse mapping from data to noise. Unlike any of the above, we develop models with hierarchical and deep latent structure, and we infer the posterior.

Recent advances in variational inference have developed expressive variational approximations (Rezende & Mohamed, 2015; Salimans et al., 2015; Tran et al., 2015). The idea of casting the design of variational families as a modeling problem was proposed in Ranganath et al. (2016b). Further advances have analyzed variational programs (Ranganath et al., 2016a)—a family of approximations which only requires a process returning samples—and which has seen further interest (Liu & Feng, 2016). Implicit-like variational approximations have also appeared in auto-encoder frameworks (Makhzani et al., 2015; Mescheder et al., 2017) and message passing (Karaletsos, 2016). We build on variational programs for inferring implicit models.

2. Implicit Probabilistic Models

In this section we describe the idea of implicit densities and develop two new classes of models: hierarchical implicit models and deep implicit models.

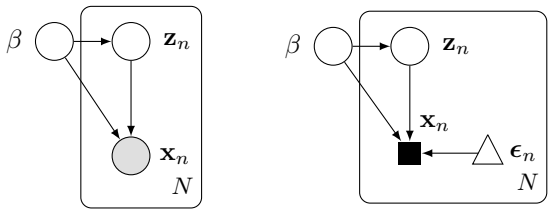

Figure 1. (**left**) Hierarchical model, with local variables $\mathbf{z}$ and global variables β. (**right**) **Hierarchical implicit model**. It is a hierarchical model where $\mathbf{x}$ is a deterministic function (denoted with a square) of noise $\boldsymbol{\epsilon}$ (denoted with a triangle).

2.1. Hierarchical Implicit Models

Hierarchical models play an important role in sharing statistical strength across examples (Gelman & Hill, 2006). The joint distribution for a broad class of hierarchical Bayesian models is

$$p(\mathbf{x}, \mathbf{z}, \boldsymbol{\beta}) = p(\boldsymbol{\beta}) \prod_{n=1}^{N} p(\mathbf{x}_n \mid \mathbf{z}_n, \boldsymbol{\beta}) p(\mathbf{z}_n \mid \boldsymbol{\beta}), \quad (1)$$

where $\mathbf{x}_n$ is an observation, $\mathbf{z}_n$ are latent variables associated to that observation (local variables), and $\boldsymbol{\beta}$ are latent variables shared across observations (global variables). Fig. 1 displays its graphical representation.

With hierarchical models, local variables can be used for clustering in mixture models, mixed memberships in topic models (Blei et al., 2003), and factors in probabilistic matrix factorization (Salakhutdinov & Mnih, 2008). Global variables can be used to pool information across data points for hierarchical regression (Gelman & Hill, 2006), topic models (Blei et al., 2003), and Bayesian nonparametrics (Teh & Jordan, 2010).

Hierarchical models typically use a tractable likelihood $p(\mathbf{x}_n \mid \mathbf{z}_n, \boldsymbol{\beta})$. But many models, such as simulator-based models (Hartig et al., 2011) and generative adversarial networks (Goodfellow et al., 2014), admit high fidelity to the true data generating process and do not admit a tractable likelihood. To overcome this limitation, we develop *hierarchical implicit models* (HIMs).

Hierarchical implicit models have the same joint factorization as Eq.1 but only assume that one can sample from the likelihood. Rather than define $p(\mathbf{x}_n \mid \mathbf{z}_n, \boldsymbol{\beta})$ explicitly, HIMs define a function g that takes in random noise $\boldsymbol{\epsilon}_n \sim s(\cdot)$ and outputs $\mathbf{x}_n$ given $\mathbf{z}_n$ and $\boldsymbol{\beta}$,

$$\mathbf{x}_n = g(\boldsymbol{\epsilon}_n \mid \mathbf{z}_n, \boldsymbol{\beta}), \quad \boldsymbol{\epsilon}_n \sim s(\cdot).$$

The induced likelihood of $\mathbf{x}_n \in A$ given $\mathbf{z}_n$ and $\boldsymbol{\beta}$ is

$$\mathcal{P}(\mathbf{x}_n \in A \mid \mathbf{z}_n, \boldsymbol{\beta}) = \int_{\{g(\boldsymbol{\epsilon}_n \mid \mathbf{z}_n, \boldsymbol{\beta}) = \mathbf{x}_n \in A\}} s(\boldsymbol{\epsilon}_n)\, \mathrm{d}\boldsymbol{\epsilon}_n.$$

This is typically intractable: finding the set to integrate over is difficult, and the integration may be expensive.

Fig. 1 displays the graphical model for HIMs. Noise ($\boldsymbol{\epsilon}_n$) are denoted by triangles; deterministic computation ($\mathbf{x}_n$) are denoted by squares. We illustrate two examples.

Example: Simulator Given initial conditions, simulators describe a stochastic process that generates data. For example, in population ecology, the Lotka-Volterra model simulates predator-prey populations over time via a stochastic differential equation (Wilkinson, 2011). For prey and predator populations $x_1, x_2 \in \mathbb{R}^+$ respectively, one process is

$$\frac{\mathrm{d}x_1}{\mathrm{d}t} = \beta_1 x_1 - \beta_2 x_1 x_2 + \epsilon_1, \qquad \epsilon_1 \sim \text{Normal}(0, 10),$$
$$\frac{\mathrm{d}x_2}{\mathrm{d}t} = -\beta_2 x_2 - \beta_3 x_1 x_2 + \epsilon_2, \quad \epsilon_2 \sim \text{Normal}(0, 10),$$

where Gaussian noise ϵ_1, ϵ_2 are added at each full time step. The simulator runs for T time steps given initial population sizes for x_1, x_2. Lognormal priors are placed over β. The Lotka-Volterra model is grounded by theory and features an intractable likelihood. We study it in Section 4.2.

Example: Bayesian Generative Adversarial Network Generative adversarial networks (GANs) define an implicit model and a method for its parameter estimation (Goodfellow et al., 2014). They are known to perform well for image generation (Radford et al., 2016). Formally, the implicit model for a generative adversarial network is

$$\mathbf{x}_n = g(\boldsymbol{\epsilon}_n; \boldsymbol{\theta}), \quad \boldsymbol{\epsilon}_n \sim s(\cdot), \tag{2}$$

where g is a neural network with parameters $\boldsymbol{\theta}$, and s is a standard normal or uniform. The neural network g may not be invertible; this makes the likelihood intractable.

We make GANs amenable to Bayesian analysis by placing a prior on the parameters $\boldsymbol{\theta}$. We call this a Bayesian GAN. Bayesian GANs enable modeling of parameter uncertainty and are inspired by Bayesian neural networks, which have been shown to improve the data efficiency of standard neural networks (MacKay, 1992; Neal, 1994). We study Bayesian GANs in Section 4.3.

Noise versus Latent Variables. With implicit models, a natural question is whether one should additionally treat the noise as latent variables and infer their posterior. The posterior's shape—and ultimately if it is meaningful—is determined by the dimensionality of noise and the transformation. For example, consider the GAN model $\mathbf{x}_n = g(\boldsymbol{\epsilon}_n; \boldsymbol{\theta})$. The conditional $p(\mathbf{x}_n \mid \boldsymbol{\epsilon}_n)$ is a point mass, fully determined by $\boldsymbol{\epsilon}_n$. When $g(\cdot; \boldsymbol{\theta})$ is injective, the posterior $p(\boldsymbol{\epsilon}_n \mid \mathbf{x}_n)$ is also a point mass,

$$p(\boldsymbol{\epsilon}_n \mid \mathbf{x}_n) = \mathbb{I}[\boldsymbol{\epsilon}_n = g^{-1}(\mathbf{x}_n)],$$

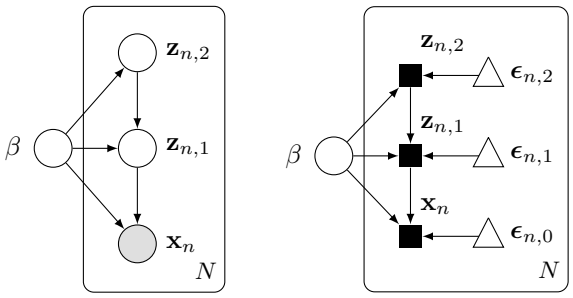

Figure 2. (**left**) Deep generative model with 2 layers $\{\mathbf{z}_1, \mathbf{z}_2\}$ and global variables $\boldsymbol{\beta}$. (**right**) **Deep implicit model**. It is a deep generative model with implicit likelihood and implicit priors.

where g^{-1} is the left inverse of g. This means for injective functions, the "posterior" may be worth analysis as a deterministic hidden representation (Donahue et al., 2017), but it is not random. (Injective g is the typical case for GANs, as the noise dimensionality is set to be lower than the data's). The point mass posterior can be found via nonlinear least squares; see Appendix A.

2.2. Deep Implicit Models

With hierarchical implicit models, we define an implicit likelihood to build complex densities. We can also define implicit latent variables. Stacking layers of implicit latent variables produces a *deep implicit model* (DIM).

DIMs share the same local-global decomposition as hierarchical implicit models. The key difference is that instead of a single local variable $\mathbf{z}_n$, DIMs define layers of implicit local variables. Formally, for layers $1, \ldots, L$, the local variables for data point $\mathbf{x}_n$ are generated via

$$\begin{aligned}
\mathbf{z}_{n,L} &= g_L(\boldsymbol{\epsilon}_{n,L} \mid \boldsymbol{\beta}_L), & \boldsymbol{\epsilon}_{n,L} &\sim s(\cdot), \\
\mathbf{z}_{n,L-1} &= g_{L-1}(\boldsymbol{\epsilon}_{n,L-1} \mid \mathbf{z}_{n,L}, \boldsymbol{\beta}_{L-1}), & \boldsymbol{\epsilon}_{n,L-1} &\sim s(\cdot), \\
&\vdots & & \\
\mathbf{z}_{n,1} &= g_1(\boldsymbol{\epsilon}_{n,1} \mid \mathbf{z}_{n,2}, \boldsymbol{\beta}_1), & \boldsymbol{\epsilon}_{n,1} &\sim s(\cdot),
\end{aligned} \tag{3}$$

where g_ℓ is the transformation for layer ℓ, s is random noise, and $\boldsymbol{\beta} = \{\boldsymbol{\beta}_1, \ldots, \boldsymbol{\beta}_L\}$ are parameters of the L transformations (e.g., weights and biases). Similar to the likelihood for $\mathbf{x}_n$, each local variable $\mathbf{z}_n$ is implicit; this enables complex prior densities.[1]

Fig. 2 displays the graphical model for a DIM. DIMs can be broadly interpreted as combining the strengths of probabilistic graphical models and generative adversarial networks. The former enables rich, interpretable latent structure (Neal, 1990; Ranganath et al., 2015) and applies to a wide range of tasks (Rezende et al., 2014; Kingma et al.,

[1] We described DIMs with the transformations' parameters as global variables $\boldsymbol{\beta}$. It is also natural to combine other globals, known as a hierarchical deep model (Salakhutdinov et al., 2013).

2014; Rezende et al., 2016). The latter enables accurate density modeling (Goodfellow et al., 2014).

DIMs can also be motivated from a computational perspective. Following Eq.3, suppose neural net layers define each g_ℓ. A DIM is a deep neural network with random noise injected at certain layers. Noise at multiple layers does not increase the model's representational capacity (one level is enough to represent all distributions). Rather, DIMs impose an inductive bias for data to exhibit multi-level structure: the noise makes the data robust to deviations from higher level concepts. This connects to the observation that training may be easier if we inject randomness at multiple layers of a neural net, rather than simply at the input (Goodfellow et al., 2016).

As an example of a DIM, suppose each layer's function g_ℓ is linear with respect to the noise,

$$g_\ell(\boldsymbol{\epsilon}_{n,\ell} \mid \mathbf{z}_{n,\ell+1}, \beta_\ell) = \mu(\mathbf{z}_{n,\ell+1}, \beta_\ell) + \sigma(\mathbf{z}_{n,\ell+1}, \beta_\ell)\boldsymbol{\epsilon}_{n,\ell}.$$

for some functions $\mu(\cdot), \sigma(\cdot)$. The induced distribution for $\mathbf{z}_{n,\ell}$ is Gaussian with parameters $\mu(\cdot)$ and $\sigma(\cdot)$. This defines a deep latent Gaussian model (Rezende et al., 2014). Nonlinear interactions with the noise $\boldsymbol{\epsilon}$, different output dimensionality, and different support of the output, expand the class of deep generative models. In Appendix B, we provide other examples, written in the Edward probabilistic programming language (Tran et al., 2016).

3. Variational Inference for Implicit Models

We described two classes of implicit models, enabling both hierarchical structure and deep latent structure. Given data, we aim to calculate the model's posterior $p(\mathbf{z}, \boldsymbol{\beta} \mid \mathbf{x}) = p(\mathbf{x}, \mathbf{z}, \boldsymbol{\beta})/p(\mathbf{x})$. This is difficult as the normalizing constant $p(\mathbf{x})$ is typically intractable. With implicit models, the lack of a likelihood function introduces an additional source of intractability.

We use variational inference (Jordan et al., 1999). It posits an approximating family $q \in \mathcal{Q}$ and optimizes to find the member closest to $p(\mathbf{z}, \boldsymbol{\beta} \mid \mathbf{x})$. There are many choices of variational objectives that measure closeness (Ranganath et al., 2016a; Li & Turner, 2016; Dieng et al., 2016). To choose an objective, we lay out desiderata for a variational inference algorithm for implicit models:

1. *Scalability*. Machine learning hinges on stochastic optimization to scale to massive data (Bottou, 2010). The variational objective should admit unbiased subsampling with the standard technique,

$$\sum_{n=1}^{N} f(\mathbf{x}_n) \approx \frac{N}{M} \sum_{m=1}^{M} f(\mathbf{x}_m),$$

where some computation $f(\cdot)$ over the full data is approximated with a mini-batch of data $\{\mathbf{x}_m\}$.

2. *Implicit Local Approximations*. Implicit models specify flexible densities; this induces very complex posterior distributions. Thus we would like a rich approximating family for the per-data point approximations $q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})$. This means the variational objective should only require that one can sample $\mathbf{z}_n \sim q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})$ and not evaluate its density.

One variational objective meeting our desiderata is based on the classical minimization of the Kullback-Leibler (KL) divergence. (Surprisingly, in Section 3.5 we outline how the KL is the *only* possible objective among a broad class.)

3.1. KL Variational Objective

The classic form of variational inference minimizes the KL divergence from the variational approximation q to the posterior distribution. This is equivalent to maximizing the *evidence lower bound* (ELBO),

$$\mathcal{L} = \mathbb{E}_{q(\boldsymbol{\beta}, \mathbf{z} \mid \mathbf{x})}[\log p(\mathbf{x}, \mathbf{z}, \boldsymbol{\beta}) - \log q(\boldsymbol{\beta}, \mathbf{z} \mid \mathbf{x})]. \quad (4)$$

Let q factorize in the same way as the posterior,

$$q(\boldsymbol{\beta}, \mathbf{z} \mid \mathbf{x}) = q(\boldsymbol{\beta}) \prod_{n=1}^{N} q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta}),$$

where $q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})$ is an intractable density and since the data $\mathbf{x}$ is constant during inference, we drop the conditioning for the global approximation $q(\boldsymbol{\beta})$.

Substituting p and q's factorization into Eq.4 yields

$$\mathcal{L} = \mathbb{E}_{q(\boldsymbol{\beta})}[\log p(\boldsymbol{\beta}) - \log q(\boldsymbol{\beta} \mid \mathbf{x})] + \sum_{n=1}^{N} \mathbb{E}_{q(\boldsymbol{\beta})q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})}[\log p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}) - \log q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})].$$

This objective presents difficulties: the local densities $p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})$ and $q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})$ are both intractable. To solve this, we consider ratio estimation.

3.2. Ratio Estimation for the KL Objective

Let $q(\mathbf{x}_n)$ be the empirical distribution on the observations $\mathbf{x}$. We subtract a constant to the objective,

$$\mathcal{L} \propto \mathbb{E}_{q(\boldsymbol{\beta})}[\log p(\boldsymbol{\beta}) - \log q(\boldsymbol{\beta})] + \sum_{n=1}^{N} \mathbb{E}_{q(\boldsymbol{\beta})q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})}[\log p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}) - \log q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})] - \log q(\mathbf{x}_n). \quad (5)$$

Denote the joint $q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}) = q(\mathbf{x}_n)q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})$. Rewrite the log differences as a ratio,

$$\mathcal{L} = \mathbb{E}_{q(\boldsymbol{\beta})}[\log p(\boldsymbol{\beta}) - \log q(\boldsymbol{\beta})] + \sum_{n=1}^{N} \mathbb{E}_{q(\boldsymbol{\beta})q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})} \left[\log \frac{p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}{q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}\right]. \tag{6}$$

Thus the ELBO is a function of the ratio of two intractable densities. If we can form an estimator of this ratio, we can proceed with optimizing the ELBO.

We apply techniques for ratio estimation (Sugiyama et al., 2012). It is a key idea in GANs (Mohamed & Lakshminarayanan, 2016; Uehara et al., 2016), and similar ideas have rearisen in statistics and physics (Gutmann et al., 2014; Cranmer et al., 2015). In particular, we use class probability estimation: given a sample from $p(\cdot)$ or $q(\cdot)$ we aim to estimate the probability that it belongs to $p(\cdot)$. We model this using $\sigma(r(\cdot; \boldsymbol{\theta}))$, where r is a parameterized function (e.g., neural network) taking sample inputs and outputting a real value; σ is the logistic function outputting the probability.

We train $r(\cdot; \boldsymbol{\theta})$ by minimizing a loss function known as a proper scoring rule (Gneiting & Raftery, 2007). For example, in experiments we use the log loss,

$$\mathcal{D}_{\text{log}} = \mathbb{E}_{p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}[-\log \sigma(r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}; \boldsymbol{\theta}))] + \mathbb{E}_{q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}[-\log(1 - \sigma(r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}; \boldsymbol{\theta})))]. \tag{7}$$

The loss is zero if $\sigma(r(\cdot; \boldsymbol{\theta}))$ returns 1 when a sample is from $p(\cdot)$ and 0 when a sample is from $q(\cdot)$. (We also experiment with the hinge loss; see Section 4.2.) If $r(\cdot; \boldsymbol{\theta})$ is sufficiently expressive, minimizing the loss obtains the optimal function,

$$r^*(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}) = \log p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}) - \log q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}).$$

This means as we minimize Eq.7, we can use our estimate $r(\cdot; \boldsymbol{\theta})$ as a proxy to the log ratio in Eq.6.[2] The gradient of $\mathcal{D}_{\text{log}}$ with respect to $\boldsymbol{\theta}$ is

$$\nabla_{\boldsymbol{\theta}} \mathcal{D}_{\text{log}} = \mathbb{E}_{p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}[\nabla_{\boldsymbol{\theta}} \log \sigma(r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}; \boldsymbol{\theta}))] + \mathbb{E}_{q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}[\nabla_{\boldsymbol{\theta}} \log(1 - \sigma(r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}; \boldsymbol{\theta})))]. \tag{8}$$

We compute unbiased gradients using Monte Carlo samples from $p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})$ and $q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})$.

3.3. Stochastic Gradients of the KL Objective

To optimize the ELBO, we use the ratio estimator,

$$\mathcal{L} = \mathbb{E}_{q(\boldsymbol{\beta} \mid \mathbf{x})}[\log p(\boldsymbol{\beta}) - \log q(\boldsymbol{\beta})] + \sum_{n=1}^{N} \mathbb{E}_{q(\boldsymbol{\beta} \mid \mathbf{x})q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})}[r(\mathbf{x}_n, \mathbf{z}_n, \beta)]. \tag{9}$$

[2]Note we estimate the log ratio. This is more numerically stable than the ratio and is of direct interest in our setting.

All terms are now tractable. We can calculate gradients to optimize the variational family q. Below we assume the priors $p(\boldsymbol{\beta}), p(\mathbf{z}_n \mid \boldsymbol{\beta})$ are differentiable. (We discuss methods to handle discrete global variables in the next section.)

We focus on reparameterizable variational approximations (Kingma & Welling, 2014; Rezende et al., 2014). They enable sampling via a differentiable transformation T of random noise, $\delta \sim s(\cdot)$. Due to Eq.9, we require the global approximation $q(\boldsymbol{\beta}; \boldsymbol{\lambda})$ to admit a tractable density. With reparameterization, its sample can be generated as

$$\boldsymbol{\beta} = T_{\text{global}}(\boldsymbol{\delta}_{\text{global}}; \boldsymbol{\lambda}), \quad \boldsymbol{\delta}_{\text{global}} \sim s(\cdot),$$

for a choice of transformation $T_{\text{global}}(\cdot; \boldsymbol{\lambda})$ and noise $s(\cdot)$. For example, setting $s(\cdot) = \mathcal{N}(0, 1)$ and $T_{\text{global}}(\boldsymbol{\delta}_{\text{global}}) = \mu + \sigma\boldsymbol{\delta}_{\text{global}}$ induces a normal distribution $\mathcal{N}(\mu, \sigma^2)$.

Similarly for the local variables $\mathbf{z}_n$, we specify

$$\mathbf{z}_n = T_{\text{local}}(\boldsymbol{\delta}_n, \mathbf{x}_n, \boldsymbol{\beta}; \boldsymbol{\phi}), \quad \boldsymbol{\delta}_n \sim s(\cdot).$$

Unlike the global approximation, the local variational density $q(\mathbf{z}_n \mid \mathbf{x}_n; \boldsymbol{\phi})$ need not be tractable: the ratio estimator relaxes this requirement. It lets us leverage implicit models not only for data but also for approximate posteriors.[3] Implicit variational families are also known as variational programs (Ranganath et al., 2016a).

The gradient with respect to global parameters $\boldsymbol{\lambda}$ under this approximating family is

$$\nabla_{\boldsymbol{\lambda}} \mathcal{L} = \mathbb{E}_{s(\boldsymbol{\delta}_{\text{global}})}[\nabla_{\boldsymbol{\lambda}}(\log p(\boldsymbol{\beta}) - \log q(\boldsymbol{\beta}))] + \sum_{n=1}^{N} \mathbb{E}_{s(\boldsymbol{\delta}_{\text{global}})s_n(\boldsymbol{\delta}_n)}[\nabla_{\boldsymbol{\lambda}} r(\mathbf{x}_n, \mathbf{z}_n, \beta)], \tag{10}$$

where the gradient backpropagates through the local sampling process $\mathbf{z}_n = T_{\text{local}}(\boldsymbol{\delta}_n, \mathbf{x}_n, \boldsymbol{\beta}; \boldsymbol{\phi})$ and the global reparameterization $\boldsymbol{\beta} = T_{\text{global}}(\boldsymbol{\delta}_{\text{global}}; \boldsymbol{\lambda})$. We compute unbiased gradients using Monte Carlo samples.

The gradient with respect to local parameters $\boldsymbol{\phi}$ is

$$\nabla_{\boldsymbol{\phi}} \mathcal{L} = \sum_{n=1}^{N} \mathbb{E}_{q(\boldsymbol{\beta})s(\boldsymbol{\delta}_n)}[\nabla_{\boldsymbol{\phi}} r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta})]. \tag{11}$$

where the gradient backpropagates through T_{local}.[4] We compute unbiased gradients using Monte Carlo.

3.4. Algorithm

Algorithm 1 outlines the procedure. We call it *likelihood-free variational inference (*LFVI*)*. LFVI is black box: it

[3]We amortize computation by sharing parameters $\boldsymbol{\phi}$ across $\mathbf{z}_n$. Following the posterior's factorization, this can be made more accurate if each local variable has its own parameters $\boldsymbol{\phi}_n$.

[4]The ratio r indirectly depends on $\boldsymbol{\phi}$ but its gradient w.r.t. $\boldsymbol{\phi}$ disappears. This is derived via the score function identity and the product rule (see, e.g., Ranganath et al. (2014, Appendix)).

Algorithm 1: Likelihood-free variational inference (LFVI)

Input : Model $\mathbf{x}_n, \mathbf{z}_n \sim p(\cdot \mid \boldsymbol{\beta}), p(\boldsymbol{\beta})$
Variational approximation
$\mathbf{z}_n \sim q(\cdot \mid \mathbf{x}_n, \boldsymbol{\beta}; \boldsymbol{\phi}), q(\boldsymbol{\beta} \mid \mathbf{x}; \boldsymbol{\lambda})$,
Ratio estimator $r(\cdot; \boldsymbol{\theta})$
Output: Variational parameters $\boldsymbol{\lambda}$, $\boldsymbol{\phi}$
Initialize $\boldsymbol{\theta}$, $\boldsymbol{\lambda}$, $\boldsymbol{\phi}$ randomly.
while *not converged* **do**
  Compute unbiased estimate of $\nabla_{\boldsymbol{\theta}} \mathcal{D}$ (Eq.8).
  Compute unbiased estimate of $\nabla_{\boldsymbol{\lambda}} \mathcal{L}$ (Eq.10).
  Compute unbiased estimate of $\nabla_{\boldsymbol{\phi}} \mathcal{L}$ (Eq.11).
  Update $\boldsymbol{\theta}$, $\boldsymbol{\lambda}$, $\boldsymbol{\phi}$ using stochastic gradient ascent.
end

applies to models in which one can simulate data and local variables, and calculate densities for the global variables. LFVI first updates $\boldsymbol{\theta}$ to improve the ratio estimator r. Then it uses r to update parameters $\{\boldsymbol{\lambda}, \boldsymbol{\phi}\}$ of the variational approximation q. We optimize r and q simultaneously. The algorithm is available in Edward (Tran et al., 2016).

LFVI is scalable: we can unbiasedly estimate the gradient over the full data set with mini-batches (Hoffman et al., 2013). The algorithm can also handle models of either continuous or discrete data. The requirement for differentiable global variables and reparameterizable global approximations can be relaxed using score function gradients (Ranganath et al., 2014).

Point estimates of the global parameters $\boldsymbol{\beta}$ suffice for many applications (Goodfellow et al., 2014; Rezende et al., 2014). Algorithm 1 can find point estimates: place a point mass approximation q on the parameters $\boldsymbol{\beta}$. This simplifies gradients and corresponds to variational EM.

3.5. The KL Uniqueness Proposition

A natural question is if there exist other objectives satisfying our desiderata of scalability and implicit local approximations. We show that the KL is special.

Proposition 1 (Uniqueness Proposition)**.** *Consider the class of integral probability metrics (Müller, 1997) and Csizar's f divergence (Csisz et al., 1967). Among these classes, the* KL *divergence is the unique divergence satisfying our desiderata.*

We sketch the argument in Appendix C.

4. Experiments

We developed new models and inference. In experiments, we first analyze the stability of the ratio estimator for training implicit models in practice. We then study three applications: a large-scale physical simulator for predator-prey populations in ecology; a Bayesian GAN for supervised classification; and a deep implicit model for text.

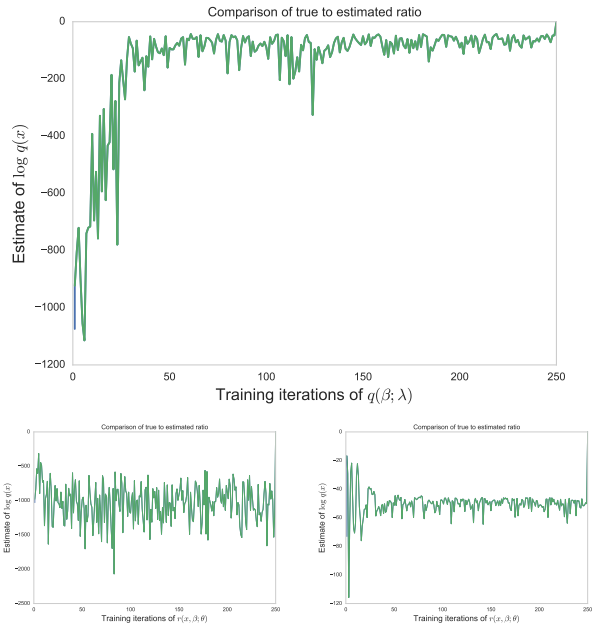
Comparison of true to estimated ratio
Estimate of $\log q(x)$
Training iterations of $q(\beta; \lambda)$
Comparison of true to estimated ratio
Training iterations of $r(x, \beta; \theta)$
Comparison of true to estimated ratio
Training iterations of $r(x, \beta; \theta)$

Figure 3. **(top)** Difference of ratios over steps of q. Low variance on y-axis means more stable. Interestingly, the ratio estimator is more accurate and stable as q converges to the posterior. **(bot. left)** Difference of ratios over steps of r; q is fixed at random initialization. The ratio estimator doesn't improve even after many steps. **(bot. right)** Difference of ratios over steps of r; q is fixed at the posterior. The ratio estimator only requires few steps from random initialization to be highly accurate.

4.1. Stability of Ratio Estimator

With implicit models, the difference from standard KL variational inference lies in the ratio estimation problem. Thus we would like to assess the accuracy of the ratio estimator. We can check this by comparing to the true ratio under a model with tractable likelihood.

We apply Bayesian linear regression. It features a tractable posterior which we leverage in our analysis. We use 50 simulated data points $\{\mathbf{x}_n \in \mathbb{R}^2, \mathbf{y}_n \in \mathbb{R}\}$. The optimal (log) ratio is

$$r^*(\mathbf{x}, \boldsymbol{\beta}) = \log p(\mathbf{x} \mid \boldsymbol{\beta}) - \log q(\mathbf{x}).$$

Note the log-likelihood $\log p(\mathbf{x} \mid \boldsymbol{\beta})$ minus $r^*(\mathbf{x}, \boldsymbol{\beta})$ is equal to the empirical distribution $\sum_n \log q(\mathbf{x}_n)$, a constant. Therefore if a ratio estimator r is accurate, its difference with $\log p(\mathbf{x} \mid \boldsymbol{\beta})$ should be a constant with low variance across values of $\boldsymbol{\beta}$.

See Fig. 3. The top graph displays the estimate of $\log q(\mathbf{x})$ over updates of the variational approximation $q(\boldsymbol{\beta})$; each

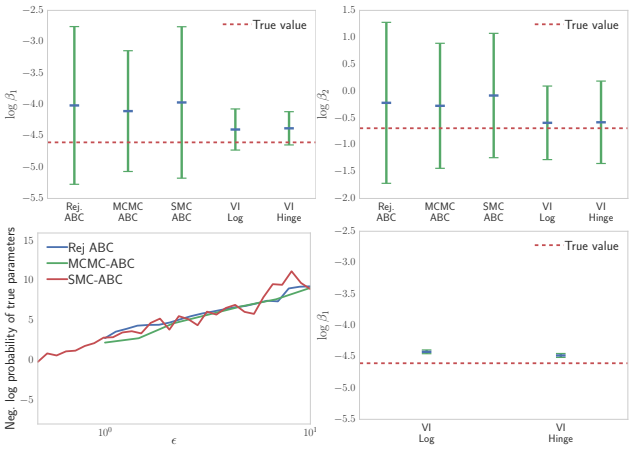

Figure 4. **(top)** Marginal posterior for first two parameters. **(bot. left)** ABC methods over tolerance error. **(bot. right)** Marginal posterior for first parameter on a large-scale data set. Our inference achieves more accurate results and scales to massive data.

estimate uses a sample from the current $q(\boldsymbol{\beta})$. The ratio estimator r is more accurate as q exactly converges to the posterior. This matches our intuition: if data generated from the model is close to the true data, then the ratio is more stable to estimate.

An alternative hypothesis for Fig. 3 is that the ratio estimator has simply accumulated information during training. This turns out to be untrue; see the bottom graphs. On the left, q is fixed at a random initialization; the estimate of $\log q(\mathbf{x})$ is displayed over updates of r. After many updates, r still produces unstable estimates. In contrast, the right shows the same procedure with q fixed at the posterior. r is accurate after few updates.

Several practical insights appear for training. First, it is not helpful to update r multiple times before updating q (at least in initial iterations). Additionally, if the specified model poorly matches the data, training will be difficult across all iterations.

The property that ratio estimation is more accurate as the variational approximation improves is because $q(\mathbf{x}_n)$ is set to be the empirical distribution. (Note we could subtract any density $q(\mathbf{x}_n)$ from the ELBO in Eq.5.) Likelihood-free variational inference finds $q(\boldsymbol{\beta})$ that makes the observed data likely under $p(\mathbf{x}_n \mid \boldsymbol{\beta})$, i.e., $p(\mathbf{x}_n \mid \boldsymbol{\beta})$ gets closer to the empirical distribution at values sampled from $q(\boldsymbol{\beta})$. Letting $q(\mathbf{x}_n)$ be the empirical distribution means the ratio estimation problem will be less trivially solvable (thus more accurate) as $q(\boldsymbol{\beta})$ improves.

4.2. Lotka-Volterra Predator-Prey Simulator

We analyze the Lotka-Volterra simulator outlined in Section 2. Its global variables $\boldsymbol{\beta}$ govern the rates of change in a simulation of predator-prey populations. To infer them, we posit a mean-field normal approximation (reparameterized to be on the same support) and run Algorithm 1 with both a log loss and hinge loss for the ratio estimation problem; also see Appendix D. We compare to rejection ABC, MCMC-ABC, and SMC-ABC (Marin et al., 2012).

Fig. 4 displays results on two data sets. In the top figures and bottom left, we analyze data consisting of a simulation for $T = 30$ time steps, with recorded values of the populations every 0.2 time units. The bottom left figure calculates the negative log probability of the true parameters over the tolerance error for ABC methods; smaller tolerances result in more accuracy but slower runtime. The top figures compare the marginal posteriors for two parameters using the smallest tolerance for the ABC methods. Rejection ABC, MCMC-ABC, and SMC-ABC all contain the true parameters in their 95% credible interval but are less confident than our methods. Further, they required $100,000$ simulations from the model, with an acceptance rate of 0.004% and 2.990% for rejection ABC and MCMC-ABC respectively.

The bottom right figure analyzes data consisting of $100,000$ time series, each of the same size as the single time series analyzed in the previous figures. This size is not possible with traditional methods. Further, we see that with our methods, the posterior concentrates near the truth. We also experienced little difference in accuracy between using the log loss or the hinge loss for ratio estimation.

4.3. Bayesian Generative Adversarial Networks

We analyze Bayesian GANs, described in Section 2. Mimicking a use case of Bayesian neural networks (Blundell et al., 2015; Hernández-Lobato et al., 2016), we apply Bayesian GANs for classification on small to medium-size data. The GAN defines a conditional $p(y_n \mid \mathbf{x}_n)$, taking a feature $\mathbf{x}_n \in \mathbb{R}^D$ as input and generating a label $y_n \in \{1, \dots, K\}$, via the process

$$y_n = g(\mathbf{x}_n, \boldsymbol{\epsilon}_n \mid \boldsymbol{\theta}), \qquad \boldsymbol{\epsilon}_n \sim \mathcal{N}(0, 1), \tag{12}$$

where $g(\cdot \mid \boldsymbol{\theta})$ is a neural net parameterized by weights and biases $\boldsymbol{\theta}$. We place normal priors, $\boldsymbol{\theta} \sim \mathcal{N}(0, 1)$.

We analyze two choices of the variational model: one with a mean-field normal approximation for $q(\boldsymbol{\theta} \mid \mathbf{x})$, and another with a point mass approximation (equivalent to maximum a posteriori). We compare to a Bayesian neural network, which uses the same generative process as Eq.12 but draws from a Categorical distribution rather than feeding noise into the neural net. We fit it separately using a mean-field normal approximation and maximum a posteriori. Table 1 shows that Bayesian GANs generally outperform their Bayesian neural net counterpart.

Note that Bayesian GANs can analyze discrete data such as

Model + Inference	Test Set Error Crabs	Pima	Covertype	MNIST
Bayesian GAN + VI	0.03	**0.232**	**0.154**	**0.0136**
Bayesian GAN + MAP	0.12	0.240	0.185	0.0283
Bayesian NN + VI	**0.02**	0.242	0.164	0.0311
Bayesian NN + MAP	0.05	0.320	0.188	0.0623

Table 1. Classification accuracy of Bayesian GAN and Bayesian neural networks across small to medium-size data sets. Bayesian GANs achieves comparable or better performance to their Bayesian neural net counterpart.

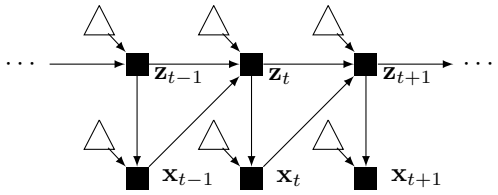

Figure 5. A deep implicit model for sequences. It is a RNN with noise injected into each hidden state. The hidden state is now an implicit latent variable. The same occurs for generating outputs.

in generating a classification label. Traditional GANs for discrete data is an open challenge (Kusner & Hernández-Lobato, 2016). In Appendix E, we compare Bayesian GANs with point estimation to typical GANs. Bayesian GANs are also able to leverage parameter uncertainty for analyzing these small to medium-size data sets.

One problem with Bayesian GANs is that they cannot work with very large neural networks: the ratio estimator is a function of global parameters, and thus the input size grows with the size of the neural network. One approach is to make the ratio estimator not a function of the global parameters. Instead of optimizing model parameters via variational EM, we can train the model parameters by backpropagating through the ratio objective instead of the variational objective. In addition, ratio estimation may be difficult in high dimensions; Appendix F discusses a potential workaround.

4.4. Deep Implicit Models

In this section, we show how to build a deep implicit model by injecting randomness into standard neural networks. We model sequences $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$. For $t = 1, \dots, T$,

$$\begin{aligned} \mathbf{z}_t &= g_z(\mathbf{x}_{t-1}, \mathbf{z}_{t-1}, \boldsymbol{\epsilon}_{t,z}), \quad & \boldsymbol{\epsilon}_{t,z} &\sim \mathcal{N}(0, 1), \\ \mathbf{x}_t &= g_x(\mathbf{z}_t, \boldsymbol{\epsilon}_{t,x}), & \boldsymbol{\epsilon}_{t,x} &\sim \mathcal{N}(0, 1), \end{aligned}$$

where g_z is a single-layer RNN and g_x serves as a decoder layer to output the data. See Fig. 5.

If the injected noise $\boldsymbol{\epsilon}_{t,z}$ combines linearly with the output of g_z, the induced distribution $p(\mathbf{z}_t \mid \mathbf{x}_{t-1}, \mathbf{z}_{t-1})$ is Gaussian parameterized by that output. This defines a stochastic RNN (Bayer & Osendorfer, 2014; Fraccaro et al., 2016) without deterministic connections. With nonlinear combinations, the implicit density is more flexible (and intractable), making previous methods for inference not applicable. In our method, we perform variational inference and specify q to be implicit; we use the same architecture as the probability model's implicit priors.

```
1 −x+x/x**x*//x*x+
2 x/x*x+x*x/x+x+x+
3 /+x*x+x*x/x/x+x+
4 /x+*x+x*x/x+x−x+
5 x/x*x/x*x+x+x+x−
6 x+x+x/x*x*x+x/x+
```

Figure 6. Generated text from the deep implicit model. Good samples place arithmetic operators between the variable x. The DIM has learned to follow rules from the context free grammar up to some multiple operator repeats.

Following Kusner & Hernández-Lobato (2016), we aim to generate simple one-variable arithmetic sequences following a context free grammar,

$$S \to x \| S + S \| S - S \| S * S \| S / S,$$

where $\|$ divides possible productions of the grammar. We concatenate the inputs and point estimate the global variables (model parameters) using variational EM. Fig. 6 displays samples from the inferred model, training on sequences with a maximum of 15 symbols. It achieves sequences which roughly follow the context free grammar.

5. Discussion: Challenges & Extensions

We developed a class of hierarchical and deep implicit models, merging the idea of implicit densities with hierarchical Bayesian modeling and deep neural networks. Further, we developed a black box variational algorithm for their accurate and scalable inference. This expands

Bayesian analysis with the ability to apply neural samplers, physical simulators, and their combination with rich, interpretable latent structure.

Stable and more automated inference are open challenges. These challenges will be especially important as we continue to build black box algorithms that make these models broadly accessible —as well as when we analyze large-scale real world applications. Further extensions also include relaxations to infer discrete latent variables, on both finite and infinite support; merging techniques for training GANs; and the design of implicit causal models.

Acknowledgements. We thank Balaji Lakshminarayanan for discussions which helped motivate this work. We also thank Christian Naesseth, Jaan Altosaar, and Adji Dieng for their feedback and comments. This work is supported by NSF IIS-1247664, ONR N00014-11-1-0651, DARPA FA8750-14-2-0009, DARPA N66001-15-C-4032, Adobe, Google, NSERC PGS-D, and the Sloan Foundation.

References

Anelli, Giovanni, Antchev, G, Aspell, P, Avati, V, Bagliesi, MG, Berardi, V, Berretti, M, Boccone, V, Bottigli, U, Bozzo, M, et al. The totem experiment at the CERN large Hadron collider. *Journal of Instrumentation*, 3(08): S08007, 2008.

Bayer, Justin and Osendorfer, Christian. Learning stochastic recurrent networks. *arXiv preprint arXiv:1411.7610*, 2014.

Beaumont, Mark A. Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution and Systematics*, 41(379-406):1, 2010.

Blei, David M, Ng, Andrew Y, and Jordan, Michael I. Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan):993–1022, 2003.

Blundell, Charles, Cornebise, Julien, Kavukcuoglu, Koray, and Wierstra, Daan. Weight uncertainty in neural network. In *International Conference on Machine Learning*, 2015.

Bottou, Léon. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pp. 177–186. Springer, 2010.

Chen, Xi, Duan, Yan, Houthooft, Rein, Schulman, John, Sutskever, Ilya, and Abbeel, Pieter. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Neural Information Processing Systems*, 2016.

Cranmer, Kyle, Pavez, Juan, and Louppe, Gilles. Approximating likelihood ratios with calibrated discriminative classifiers. *arXiv preprint arXiv:1506.02169*, 2015.

Csisz, I et al. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.

Diaconis, P, Holmes, S, and Montgomery, R. Dynamical bias in the coin toss. *SIAM*, 49(2):211–235, 2007.

Dieng, Adji B, Tran, Dustin, Ranganath, Rajesh, Paisley, John, and Blei, David M. The χ-Divergence for Approximate Inference. *arXiv preprint arXiv:1611.00328*, 2016.

Diggle, Peter J and Gratton, Richard J. Monte Carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society: Series B (Methodological)*, pp. 193–227, 1984.

Donahue, Jeff, Krähenbühl, Philipp, and Darrell, Trevor. Adversarial feature learning. In *International Conference on Learning Representations*, 2017.

Dumoulin, Vincent, Belghazi, Ishmael, Poole, Ben, Lamb, Alex, Arjovsky, Martin, Mastropietro, Olivier, and Courville, Aaron. Adversarially learned inference. In *International Conference on Learning Representations*, 2017.

Dziugaite, Gintare Karolina, Roy, Daniel M, and Ghahramani, Zoubin. Training generative neural networks via maximum mean discrepancy optimization. In *Uncertainty in Artificial Intelligence*, 2015.

Fisher, Ronald A. Inverse probability. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 26, pp. 528–535, 1930.

Fraccaro, Marco, Sønderby, Søren Kaae, Paquet, Ulrich, and Winther, Ole. Sequential neural models with stochastic layers. In *Neural Information Processing Systems*, 2016.

Gelman, Andrew and Hill, Jennifer. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press, 2006.

Gelman, Andrew, Carlin, John B, Stern, Hal S, Dunson, David B, Vehtari, Aki, and Rubin, Donald B. *Bayesian data analysis*. Texts in Statistical Science Series. CRC Press, Boca Raton, FL, 2013.

Gneiting, Tilmann and Raftery, Adrian E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

Goodfellow, I, Pouget-Abadie, J, Mirza, M, Xu, Bing, Warde-Farley, David, Ozair, Sherjil, Courville, Aaron,

and Bengio, Yoshua. Generative adversarial nets. In *Neural Information Processing Systems*, 2014.

Goodfellow, Ian, Bengio, Y, and Courville, A. *Deep learning*. MIT Press, 2016.

Goodfellow, Ian J. On distinguishability criteria for estimating generative models. In *ICLR Workshop*, 2014.

Gutmann, Michael U, Dutta, Ritabrata, Kaski, Samuel, and Corander, Jukka. Statistical Inference of Intractable Generative Models via Classification. *arXiv preprint arXiv:1407.4981*, 2014.

Hartig, Florian, Calabrese, Justin M, Reineking, Björn, Wiegand, Thorsten, and Huth, Andreas. Statistical inference for stochastic simulation models–theory and application. *Ecology Letters*, 14(8):816–827, 2011.

Hernández-Lobato, José Miguel, Li, Yingzhen, Rowland, Mark, Hernández-Lobato, Daniel, Bui, Thang, and Turner, Richard E. Black-box α-divergence minimization. In *International Conference on Machine Learning*, 2016.

Hoffman, Matthew D, Blei, David M, Wang, Chong, and Paisley, John William. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.

Jordan, Michael I, Ghahramani, Zoubin, Jaakkola, Tommi S, and Saul, Lawrence K. An introduction to variational methods for graphical models. *Machine Learning*, 1999.

Karaletsos, Theofanis. Adversarial message passing for graphical models. In *NIPS Workshop*, 2016.

Keller, Joseph B. The probability of heads. *The American Mathematical Monthly*, 93(3):191–197, 1986.

Kingma, Diederik P and Welling, Max. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.

Kingma, Diederik P, Rezende, Danilo J, Mohamed, Shakir, and Welling, Max. Semi-supervised learning with deep generative models. In *Neural Information Processing Systems*, 2014.

Kusner, Matt J and Hernández-Lobato, José Miguel. GANs for sequences of discrete elements with the Gumbel-Softmax distribution. In *NIPS Workshop*, 2016.

Laplace, PS. A philosophical essay on probabilities. 1814.

Larsen, Anders Boesen Lindbo, Sønderby, Søren Kaae, Larochelle, Hugo, and Winther, Ole. Autoencoding beyond pixels using a learned similarity metric. In *International Conference on Machine Learning*, 2016.

Li, Yingzhen and Turner, Richard E. Rényi Divergence Variational Inference. In *Neural Information Processing Systems*, 2016.

Liu, Qiang and Feng, Yihao. Two methods for wild variational inference. *arXiv preprint arXiv:1612.00081*, 2016.

MacKay, David J C. *Bayesian methods for adaptive models*. PhD thesis, California Institute of Technology, 1992.

Makhzani, Alireza, Shlens, Jonathon, Jaitly, Navdeep, and Goodfellow, Ian. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.

Marin, Jean-Michel, Pudlo, Pierre, Robert, Christian P, and Ryder, Robin J. Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180, 2012.

Mescheder, Lars, Nowozin, Sebastian, and Geiger, Andreas. Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. *arXiv preprint arXiv:1701.04722*, 2017.

Mohamed, Shakir and Lakshminarayanan, Balaji. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.

Müller, Alfred. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, pp. 429–443, 1997.

Murphy, Kevin. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.

Neal, Radford M. Learning stochastic feedforward networks. Technical report, 1990.

Neal, Radford M. *Bayesian Learning for Neural Networks*. PhD thesis, University of Toronto, 1994.

Pearl, Judea. *Causality*. Cambridge University Press, 2000.

Pritchard, J K, Seielstad, M T, Perez-Lezaun, A, and Feldman, M W. Population growth of human Y chromosomes: a study of Y chromosome microsatellites. *Molecular Biology and Evolution*, 16(12):1791–1798, 1999.

Radford, Alec, Metz, Luke, and Chintala, Soumith. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations*, 2016.

Ranganath, Rajesh, Gerrish, Sean, and Blei, David M. Black box variational inference. In *Artificial Intelligence and Statistics*, 2014.

Ranganath, Rajesh, Tang, Linpeng, Charlin, Laurent, and Blei, David M. Deep exponential families. In *Artificial Intelligence and Statistics*, 2015.

Ranganath, Rajesh, Altosaar, Jaan, Tran, Dustin, and Blei, David M. Operator variational inference. In *Neural Information Processing Systems*, 2016a.

Ranganath, Rajesh, Tran, Dustin, and Blei, David M. Hierarchical variational models. In *International Conference on Machine Learning*, 2016b.

Rezende, Danilo J and Mohamed, Shakir. Variational inference with normalizing flows. In *International Conference on Machine Learning*, 2015.

Rezende, Danilo J, Mohamed, Shakir, and Wierstra, Daan. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, 2014.

Rezende, Danilo Jimenez, Mohamed, Shakir, Danihelka, Ivo, Gregor, Karol, and Wierstra, Daan. One-shot generalization in deep generative models. In *International Conference on Machine Learning*, 2016.

Salakhutdinov, R, Tenenbaum, J B, and Torralba, A. Learning with hierarchical-deep models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8): 1958–1971, 2013.

Salakhutdinov, Ruslan and Mnih, Andriy. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *International Conference on Machine Learning*, pp. 880–887. ACM, 2008.

Salimans, Tim, Kingma, Diederik P, and Welling, Max. Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, 2015.

Sugiyama, M, Suzuki, T, and Kanamori, T. Density-ratio matching under the Bregman divergence: A unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 2012.

Teh, Yee Whye and Jordan, Michael I. Hierarchical Bayesian nonparametric models with applications. *Bayesian Nonparametrics*, 1, 2010.

Tran, Dustin, Blei, David M, and Airoldi, Edoardo M. Copula variational inference. In *Neural Information Processing Systems*, 2015.

Tran, Dustin, Kucukelbir, Alp, Dieng, Adji B., Rudolph, Maja, Liang, Dawen, and Blei, David M. Edward: A library for probabilistic modeling, inference, and criticism. *arXiv preprint arXiv:1610.09787*, 2016.

Uehara, Masatoshi, Sato, Issei, Suzuki, Masahiro, Nakayama, Kotaro, and Matsuo, Yutaka. Generative adversarial nets from a density ratio estimation perspective. *arXiv preprint arXiv:1610.02920*, 2016.

Wilkinson, Darren J. *Stochastic modelling for systems biology*. CRC press, 2011.

A. Noise versus Latent Variables

Nonlinear least squares yields the iterative algorithm

$$\hat{\boldsymbol{\epsilon}}_n = \hat{\boldsymbol{\epsilon}}_n - \rho_t \nabla_{\hat{\boldsymbol{\epsilon}}_n} f(\hat{\boldsymbol{\epsilon}}_n)^\top (f(\hat{\boldsymbol{\epsilon}}_n) - \mathbf{x}_n),$$

for some step size sequence ρ_t. Note the updates will get stuck when the gradient of f is zero. However, the injective property of f allows the iteration to be checked for correctness (simply check if $f(\hat{\boldsymbol{\epsilon}}_n) = \mathbf{x}_n$).

B. Implicit Model Examples in Edward

We demonstrate the simplicity of implicit models via example implementations in Edward (Tran et al., 2016).

Fig. 7 implements a 2-layer deep implicit model. It uses Keras to define neural networks: `Dense(256)(x)` applies a fully connected layer with 256 hidden units and input x; weight and bias parameters are abstracted from the user. The program generates N data points $\mathbf{x}_n \in \mathbb{R}^{10}$ using two layers of implicit latent variables $\mathbf{z}_{n,1}, \mathbf{z}_{n,2} \in \mathbb{R}^d$ and with an implicit likelihood.

Fig. 8 implements a Bayesian GAN for classification. It manually defines a 2-layer neural network, where for each data index, it takes features $\mathbf{x}_n \in \mathbb{R}^{500}$ concatenated with noise $\boldsymbol{\epsilon}_n \in \mathbb{R}$ as input. The output is a label $\mathbf{y}_n \in \{-1, 1\}$, given by the sign of the last layer. We place a standard normal prior over all weights and biases. Running this program while feeding the placeholder $\mathbf{X} \in \mathbb{R}^{N \times 500}$ generates a vector of labels $\mathbf{y} \in \{-1, 1\}^N$.

```
import tensorflow as tf
from edward.models import Normal
from keras.layers import Dense

# random noise is Normal(0, 1)
eps2 = Normal(tf.zeros([N, d]), tf.ones([N, d]))
eps1 = Normal(tf.zeros([N, d]), tf.ones([N, d]))
eps0 = Normal(tf.zeros([N, d]), tf.ones([N, d]))

# alternate latent layers z with hidden layers h
z2 = Dense(128, activation='relu')(eps2)
h2 = Dense(128, activation='relu')(z2)
z1 = Dense(128, activation='relu')(tf.concat([eps1, h2], 1))
h1 = Dense(128, activation='relu')(z1)
x  = Dense(10, activation=None)(tf.concat([eps0, h1], 1))
```

Figure 7. Two-layer deep implicit model for data points $\mathbf{x}_n \in \mathbb{R}^{10}$. The architecture alternates with stochastic and deterministic layers. To define a stochastic layer, we simply inject noise by concatenating it into the input of a neural net layer.

```
import tensorflow as tf
from edward.models import Normal

# weights and biases have Normal(0, 1) prior
W1 = Normal(tf.zeros([501, 100]), tf.ones([501, 100]))
W2 = Normal(tf.zeros([100,   1]), tf.ones([100,   1]))
b1 = Normal(tf.zeros(100), tf.ones(100))
b2 = Normal(tf.zeros(1),   tf.ones(1))

# set up inputs to neural network
X = tf.placeholder(tf.float32, [N, 500])
eps = Normal(tf.zeros([N, 1]), tf.ones([N, 1]))

# y = neural_network([x, eps])
input = tf.concat([X, eps], 1)
h1 = tf.nn.relu(tf.matmul(input, W1) + b1)
h2 = tf.matmul(h1, W2) + b2
y = tf.reshape(tf.sign(h2), [-1])  # take sign, then flatten
```

Figure 8. Bayesian GAN for classification, taking $\mathbf{X} \in \mathbb{R}^{N \times 500}$ as input and generating a vector of labels $\mathbf{y} \in \{-1, 1\}^N$. The neural network directly generates the data rather than parameterizing a probability distribution.

C. KL Uniqueness Proposition Sketch

An integral probability metric measures distance between two distributions p and q,

$$d(p, q) = \sup_{f \in \mathcal{F}} |\mathbb{E}_p f - \mathbb{E}_q f|.$$

Integral probability metrics have been used for parameter estimation in generative models (Dziugaite et al., 2015) and for variational inference in models with tractable density (Ranganath et al., 2016b). In contrast to models with only local latent variables, to infer the posterior, we need an integral probability metric between it and the variational approximation. The direct approach fails because sampling from the posterior is intractable.

An indirect approach requires constructing a sufficiently broad class of functions with posterior expectation zero based on Stein's method (Ranganath et al., 2016b). These constructions require a likelihood function and its gradient. Working around the likelihood would require a form of nonparametric density estimation; unlike ratio estimation, we are unaware of a solution that sufficiently scales to high dimensions.

As another class of divergences, the f divergence is

$$d(p, q) = \mathbb{E}_q \Big[f\left(\frac{p}{q}\right) \Big].$$

Unlike integral probability metrics, f divergences are naturally conducive to ratio estimation, enabling implicit p and implicit q. However, the challenge lies in scalable computation. To subsample data in hierarchical models, we need f to satisfy up to constants $f(ab) = f(a) + f(b)$, so that the expectation becomes a sum over individual data points. For continuous functions, this is a defining property of the

log function. This implies the KL-divergence from q to p is the only f divergence where the subsampling technique in our desiderata is possible.

D. Hinge Loss

Let $r(\mathbf{x}_i, \mathbf{z}_i, \boldsymbol{\beta}; \theta)$ output a real value, as with the log loss in Section 3. The hinge loss is

$$\begin{aligned}\mathcal{D}_{\text{hinge}} = &\mathbb{E}_{p(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}[\max(0, 1 - r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}; \boldsymbol{\theta}))] + \\ &\mathbb{E}_{q(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta})}[\max(0, 1 + r(\mathbf{x}_n, \mathbf{z}_n, \boldsymbol{\beta}; \boldsymbol{\theta}))].\end{aligned}$$

We minimize this loss function by following unbiased gradients. The gradients are calculated analogously as for the log loss. The optimal r^* is the log ratio.

E. Comparing Bayesian GANs with MAP to GANs with MLE

In Section 3, we argued that MAP estimation with a Bayesian GAN enables analysis over discrete data, but GANs—even with a maximum likelihood objective (Goodfellow, 2014)—cannot. This is a surprising result: assuming a flat prior for MAP, the two are ultimately optimizing the same objective. We compare the two below.

For GANs, assume the discriminator outputs a logit probability, so that it's unconstrained instead of on $[0, 1]$. GANs with MLE use the discriminative problem

$$\max_{\boldsymbol{\theta}} \mathbb{E}_{q(\mathbf{x})}[\log \sigma(D(\mathbf{x}; \boldsymbol{\theta}))] + \mathbb{E}_{p(\mathbf{x}; \mathbf{w})}[\log(1 - \sigma(D(\mathbf{x}; \boldsymbol{\theta})))].$$

They use the generative problem

$$\min_{\mathbf{w}} \mathbb{E}_{p(\mathbf{x}; \mathbf{w})}[-\exp(D(\mathbf{x}))].$$

Solving the generative problem with reparameterization gradients requires backpropagating through data generated from the model, $\mathbf{x} \sim p(\mathbf{x}; \mathbf{w})$. This is not possible for discrete $\mathbf{x}$. Further, the exponentiation also makes this objective numerically unstable and thus unusable in practice.

Contrast this with Bayesian GANs with MLE (MAP and a flat prior). This applies a point mass variational approximation $q(\mathbf{w}') = \mathbb{I}[\mathbf{w}' = \mathbf{w}]$. It maximizes the ELBO,

$$\max_{\mathbf{w}} \mathbb{E}_{q(\mathbf{w})}[\log p(\mathbf{w}) - \log q(\mathbf{w})] + \sum_{n=1}^{N} r(\mathbf{x}_n, \mathbf{w}).$$

The first term is zero for a flat prior $p(\mathbf{w}) \propto 1$ and point mass approximation; the problem reduces to

$$\max_{\mathbf{w}} \sum_{n=1}^{N} r(\mathbf{x}_n, \mathbf{w}).$$

Solving this is possible for discrete $\mathbf{x}$: it only requires backpropagating gradients through $r(\mathbf{x}, \mathbf{w})$ with respect to $\mathbf{w}$, all of which is differentiable. Further, the objective does not require a numerically unstable exponentiation.

Ultimately, the difference lies in the role of the ratio estimators. Recall for Bayesian GANs, we use the ratio estimation problem

$$\begin{aligned}\mathcal{D}_{\log} = &\mathbb{E}_{p(\mathbf{x}; \mathbf{w})}[-\log \sigma(r(\mathbf{x}, \mathbf{w}; \boldsymbol{\theta}))] + \\ &\mathbb{E}_{q(\mathbf{x})}[-\log(1 - \sigma(r(\mathbf{x}, \mathbf{w}; \boldsymbol{\theta})))].\end{aligned}$$

The optimal ratio estimator is the log-ratio $r^*(\mathbf{x}, \mathbf{w}) = \log p(\mathbf{x} \mid \mathbf{w}) - \log q(\mathbf{x})$. Optimizing it with respect to $\mathbf{w}$ reduces to optimizing the log-likelihood $\log p(\mathbf{x} \mid \mathbf{w})$. The optimal discriminator for GANs with MLE has the same ratio, $D^*(\mathbf{x}) = \log p(\mathbf{x}; \mathbf{w}) - \log q(\mathbf{x})$; however, it is a constant function with respect to $\mathbf{w}$. Hence one cannot immediately substitute $D^*(\mathbf{x})$ as a proxy to optimizing the likelihood. An alternative is to use importance sampling; the result is the former objective (Goodfellow, 2014).

F. Splitting the Ratio

High dimensional distributions separate more easily. This means they are easier to classify, and ratio estimation becomes unstable. One way to mitigate this problem is to split the ratio into ratios for each variable. Formally, we can estimate the ratios

$$\begin{aligned}r_{\mathbf{x}}(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}) &= \log \frac{p(\mathbf{x}_n \mid \mathbf{z}_n, \boldsymbol{\beta})}{q(\mathbf{x}_n)}, \\ r_{\mathbf{z}}(\mathbf{x}_n, \mathbf{z}_n \mid \boldsymbol{\beta}) &= \log \frac{p(\mathbf{z}_n \mid \boldsymbol{\beta})}{q(\mathbf{z}_n \mid \mathbf{x}_n, \boldsymbol{\beta})}.\end{aligned}$$

These add to form the full log ratio for Eq.9.